Extensive Form Games

CPSC 532A Lecture 7

Lecture Overview

1. Recap
2. Computing Correlated Equilibria
3. Perfect-Information Extensive-Form Games
4. Subgame Perfection

Computational Problems in Domination

- Identifying strategies dominated by a pure strategy
  - polynomial, straightforward algorithm
- Identifying strategies dominated by a mixed strategy
  - polynomial, somewhat tricky LP
- Identifying strategies that survive iterated elimination
  - repeated calls to the above LP
- Asking whether a strategy survives iterated elimination under all elimination orderings
  - polynomial for strict domination (elimination doesn't matter)
  - NP-complete otherwise

Rationalizability

- Rather than ask what is irrational, ask what is a best response to some beliefs about the opponent
  - assumes opponent is rational
  - assumes opponent knows that you and the others are rational
  - ...
- Equilibrium strategies are always rationalizable; so are lots of other strategies (but not everything).
- In two-player games, rationalizable $\Leftrightarrow$ survives iterated removal of strictly dominated strategies.

Formal definition

Definition (Correlated equilibrium)

Given an n-agent game $G = (N, A, u)$, a correlated equilibrium is a tuple (v, π, σ), where v is a tuple of random variables $v = (v_1, \dots, v_n)$ with respective domains $D = (D_1, \dots, D_n)$, π is a joint distribution over v, $\sigma = (\sigma_1, \dots, \sigma_n)$ is a vector of mappings $\sigma_i : D_i \mapsto A_i$, and for each agent i and every mapping $\sigma'_i : D_i \mapsto A_i$ it is the case that

$$\sum_{d \in D} \pi(d) u_i \left(\sigma_1(d_1), \dots, \sigma_n(d_n)\right) \geq \sum_{d \in D} \pi(d) u_i \left(\sigma'_1(d_1), \dots, \sigma'_n(d_n)\right).$$

Existence

Theorem

For every Nash equilibrium σ^* there exists a corresponding correlated equilibrium σ.

- This is easy to show:
  - let $D_i = A_i$
  - let $\pi(d) = \prod_{i \in N} \sigma_i^*(d_i)$
  - σ_i maps each d_i to the corresponding a_i.
- Thus, correlated equilibria always exist

Remarks

- Not every correlated equilibrium is equivalent to a Nash equilibrium
  - thus, correlated equilibrium is a weaker notion than Nash
- Any convex combination of the payoffs achievable under correlated equilibria is itself realizable under a correlated equilibrium
  - start with the Nash equilibria (each of which is a CE)
  - introduce a second randomizing device that selects which CE the agents will play
  - regardless of the probabilities, no agent has incentive to deviate
  - the probabilities can be adjusted to achieve any convex combination of the equilibrium payoffs
  - the randomizing devices can be combined

Lecture Overview

1. Recap
2. Computing Correlated Equilibria
3. Perfect-Information Extensive-Form Games
4. Subgame Perfection

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a'_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i' \in a} p(a) u_i(a_i', a_{-i}) \quad \forall i \in N, \forall a_i, a_i' \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$
- we could find the social-welfare maximizing CE by adding an objective function

$$\text{maximize:} \quad \sum_{a \in A} p(a) \sum_{i \in N} u_i(a).$$

Why are CE easier to compute than NE?

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a'_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- intuitively, correlated equilibrium has only a single randomization over outcomes, whereas in NE this is constructed as a product of independent probabilities.
- To change this program so that it finds NE, the first constraint would be

$$\sum_{a \in A} u_i(a) \prod_{j \in N} p_j(a_j) \geq \sum_{a \in A} u_i(a'_i, a_{-i}) \prod_{j \in N \setminus \{i\}} p_j(a_j) \quad \forall i \in N, \forall a'_i \in A_i.$$

- This is a nonlinear constraint!

Lecture Overview

Introduction

- The normal form game representation does not incorporate any notion of sequence, or time, of the actions of the players
- The extensive form is an alternative representation that makes the temporal structure explicit.
- Two variants:
  - perfect information extensive-form games
  - imperfect-information extensive-form games

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N is a set of n players

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A is a (single) set of actions

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A
- Choice nodes and labels for these nodes:
  - Choice nodes: H is a set of non-terminal choice nodes

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A
- Choice nodes and labels for these nodes:
  - Choice nodes: H
  - Action function: $\chi : H \rightarrow 2^A$ assigns to each choice node a set of possible actions

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A
- Choice nodes and labels for these nodes:
  - Choice nodes: H
  - Action function: $\chi : H \rightarrow 2^A$
  - Player function: $\rho : H \rightarrow N$ assigns to each non-terminal node h a player $i \in N$ who chooses an action at h

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A
- Choice nodes and labels for these nodes:
  - Choice nodes: H
  - Action function: $\chi : H \rightarrow 2^A$
  - Player function: $\rho : H \rightarrow N$
- Terminal nodes: Z is a set of terminal nodes, disjoint from H

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A
- Choice nodes and labels for these nodes:
  - Choice nodes: H
  - Action function: $\chi : H \rightarrow 2^A$
  - Player function: $\rho : H \rightarrow N$
- Terminal nodes: Z
- Successor function: $\sigma : H \times A \rightarrow H \cup Z$ maps a choice node and an action to a new choice node or terminal node such that for all $h_1, h_2 \in H$ and $a_1, a_2 \in A$, if $\sigma(h_1, a_1) = \sigma(h_2, a_2)$ then $h_1 = h_2$ and $a_1 = a_2$
  - The choice nodes form a tree, so we can identify a node with its history.

Definition

A (finite) perfect-information game (in extensive form) is defined by the tuple $(N, A, H, Z, \chi, \rho, \sigma, u)$, where:

- Players: N
- Actions: A
- Choice nodes and labels for these nodes:
  - Choice nodes: H
  - Action function: $\chi : H \rightarrow 2^A$
  - Player function: $\rho : H \rightarrow N$
- Terminal nodes: Z
- Successor function: $\sigma : H \times A \rightarrow H \cup Z$
- Utility function: $u = (u_1, \ldots, u_n)$; $u_i : Z \rightarrow \mathbb{R}$ is a utility function for player i on the terminal nodes Z

Example: the sharing game

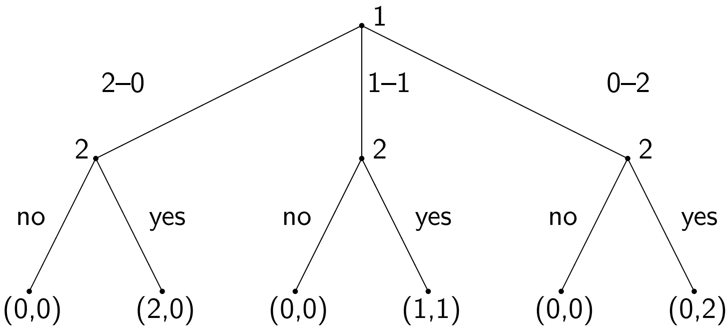

Example: the sharing game

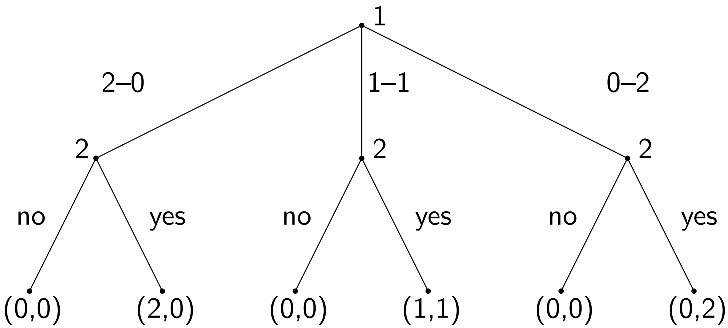

Play as a fun game, dividing 100 dollar coins. (Play each partner only once.)

Pure Strategies

- In the sharing game (splitting 2 coins) how many pure strategies does each player have?

Pure Strategies

- In the sharing game (splitting 2 coins) how many pure strategies does each player have?
  - player 1: 3; player 2: 8

Pure Strategies

- In the sharing game (splitting 2 coins) how many pure strategies does each player have?
  - player 1: 3; player 2: 8
- Overall, a pure strategy for a player in a perfect-information game is a complete specification of which deterministic action to take at every node belonging to that player.

Definition (pure strategies)

Let $G = (N, A, H, Z, \chi, \rho, \sigma, u)$ be a perfect-information extensive-form game. Then the pure strategies of player i consist of the cross product

$$\bigtimes_{h \in H, \rho(h)=i} \chi(h)$$

Pure Strategies Example

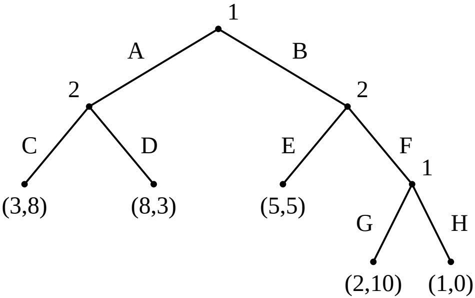

What are the pure strategies for player 2?

Pure Strategies Example

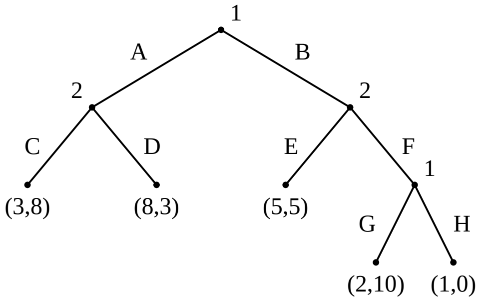

What are the pure strategies for player 2?

- $S_2 = \{(C, E); (C, F); (D, E); (D, F)\}$

Pure Strategies Example

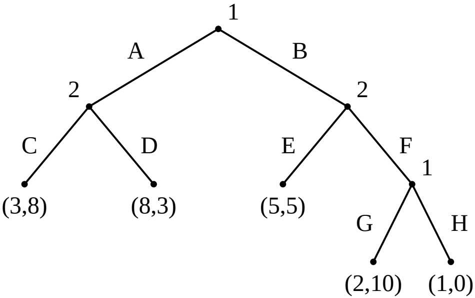

What are the pure strategies for player 2?

- $S_2 = \{(C, E); (C, F); (D, E); (D, F)\}$

What are the pure strategies for player 1?

Pure Strategies Example

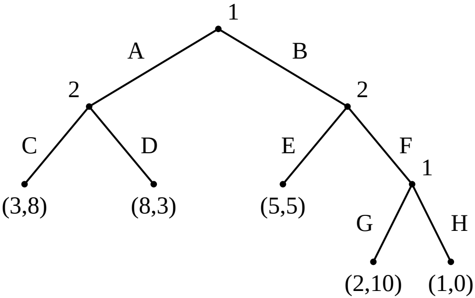

What are the pure strategies for player 2?

- $S_2 = \{(C, E); (C, F); (D, E); (D, F)\}$

What are the pure strategies for player 1?

- $S_1 = \{(B, G); (B, H), (A, G), (A, H)\}$
- This is true even though, conditional on taking A, the choice between G and H will never have to be made

Nash Equilibria

Given our new definition of pure strategy, we are able to reuse our old definitions of:

- mixed strategies
- best response
- Nash equilibrium

Theorem

Every perfect information game in extensive form has a PSNE

This is easy to see, since the players move sequentially.

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

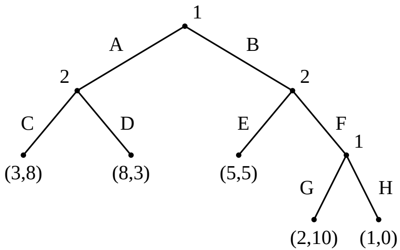

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

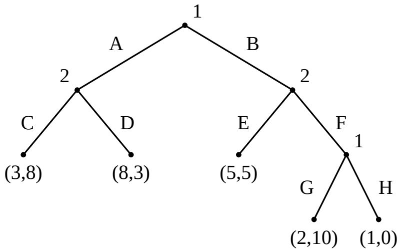

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

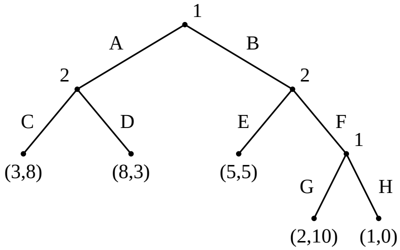

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

- this illustrates the lack of compactness of the normal form
  - games aren't always this small
  - even here we write down 16 payoff pairs instead of 5

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

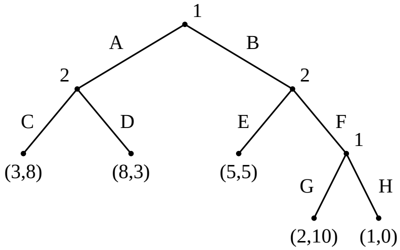

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

- while we can write any extensive-form game as a NF, we can't do the reverse.
  - e.g., matching pennies cannot be written as a perfect-information extensive form game

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

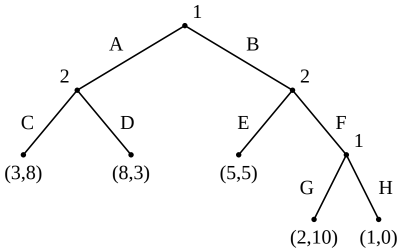

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

- What are the (three) pure-strategy equilibria?

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

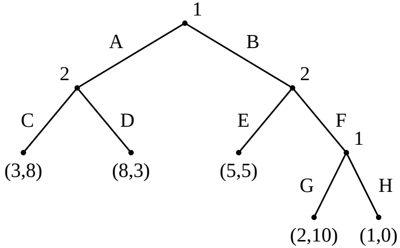

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

- What are the (three) pure-strategy equilibria?
  - $(A,G),(C,F)$
  - $(A,H),(C,F)$
  - $(B,H),(C,E)$

Induced Normal Form

- In fact, the connection to the normal form is even tighter
  - we can "convert" an extensive-form game into normal form

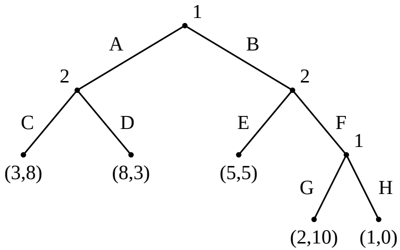

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

- What are the (three) pure-strategy equilibria?
  - $(A,G),(C,F)$
  - $(A,H),(C,F)$
  - $(B,H),(C,E)$

Lecture Overview

1. Recap
2. Computing Correlated Equilibria
3. Perfect-Information Extensive-Form Games
4. Subgame Perfection

Subgame Perfection

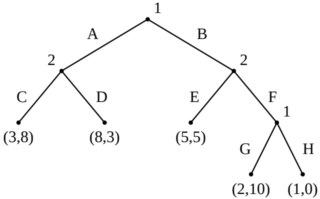

- There's something intuitively wrong with the equilibrium $(B, H), (C, E)$
  - Why would player 1 ever choose to play H if he got to the second choice node?
    - After all, G dominates H for him

Subgame Perfection

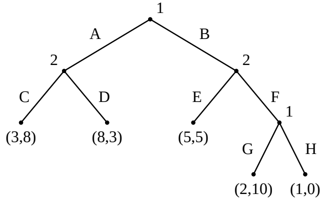

- There's something intuitively wrong with the equilibrium $(B, H), (C, E)$
  - Why would player 1 ever choose to play H if he got to the second choice node?
    - After all, G dominates H for him
  - He does it to threaten player 2, to prevent him from choosing F, and so gets 5
    - However, this seems like a non-credible threat
    - If player 1 reached his second decision node, would he really follow through and play H?

Formal Definition

- Define subgame of G rooted at h:
  - the restriction of G to the descendents of H.
- Define set of subgames of G:
  - subgames of G rooted at nodes in G

- s is a subgame perfect equilibrium of G iff for any subgame G' of G, the restriction of s to G' is a Nash equilibrium of G'
- Notes:
  - since G is its own subgame, every SPE is a NE.
  - this definition rules out "non-credible threats"

Which equilibria are subgame perfect?

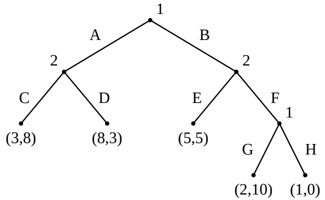

- Which equilibria from the example are subgame perfect?
  - $(A, G), (C, F)$:
  - $(B, H), (C, E)$:
  - $(A, H), (C, F)$:

Which equilibria are subgame perfect?

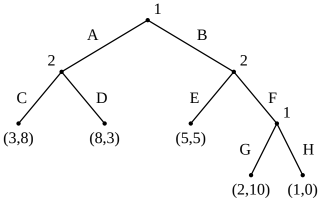

- Which equilibria from the example are subgame perfect?
  - $(A, G), (C, F)$: is subgame perfect
  - $(B, H), (C, E)$:
  - $(A, H), (C, F)$:

Which equilibria are subgame perfect?

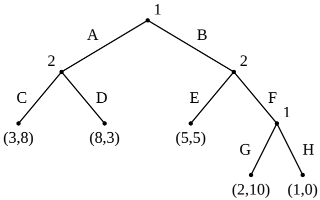

- Which equilibria from the example are subgame perfect?
  - $(A, G), (C, F)$: is subgame perfect
  - $(B, H), (C, E)$: (B, H) is an non-credible threat; not subgame perfect
  - $(A, H), (C, F)$:

Which equilibria are subgame perfect?

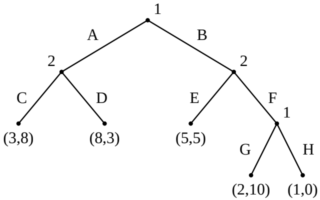

- Which equilibria from the example are subgame perfect?
  - $(A, G), (C, F)$: is subgame perfect
  - $(B, H), (C, E)$: (B, H) is an non-credible threat; not subgame perfect
  - $(A, H), (C, F)$: (A, H) is also non-credible, even though H is "off-path"